

Temporal Action Detection with Structured Segment Networks

Yue Zhao^{*1}, Yuanjun Xiong^{*1}, Limin Wang², Zhirong Wu¹, Dahua Lin¹, and Xiaoou Tang¹

¹Department of Information Engineering, The Chinese University of Hong Kong

²Computer Vision Laboratory, ETH Zurich, Switzerland

Abstract

Detecting activities in untrimmed videos is an important yet challenging task. In this paper, we tackle the difficulties of effectively locating the start and the end of a long complex action, which are often met by existing methods. Our key contribution is the structured segment network, a novel framework for temporal action detection, which models the temporal structure of each activity instance via a structured temporal pyramid. On top of the pyramid, we further introduce a decomposed discriminative model, which comprises two classifiers, respectively for classifying activities and determining completeness. This allows the framework to effectively distinguish positive proposals from background or incomplete ones, thus leading to both accurate recognition and localization. These components are integrated into a unified network that can be efficiently trained in an end-to-end fashion. We also propose a simple yet effective temporal action proposal scheme that can generate proposals of considerably higher qualities. On two challenging benchmarks, THUMOS14 and ActivityNet, our method remarkably outperforms existing state-of-the-art methods by over 10% absolute average mAP, demonstrating superior accuracy and strong adaptivity in handling activities with various temporal structures. ¹

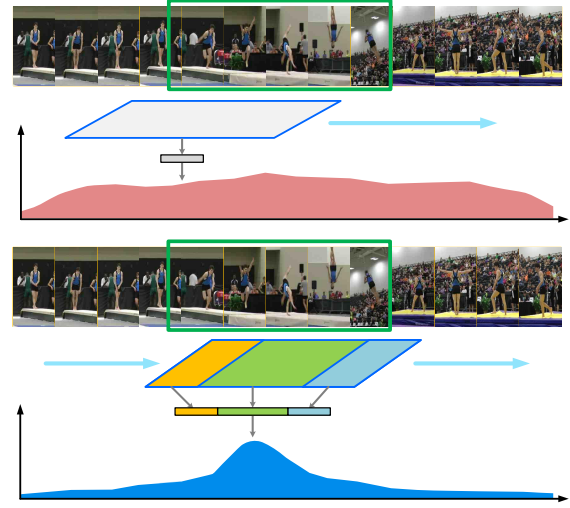


Figure 1. Importance of modeling stage structures in action detection. We slide window detectors through a video clip with an activity instance of “Tumbling” (green box). **Top:** The detector builds features without any stage structure of the activity, e.g. average pooling throughout the window. It produces high responses everywhere, making it hard to localize the instance. **Bottom:** Our detector utilizes stage structures (starting, course, and ending) via structured temporal pyramid pooling. Its response is only significant when the window is well aligned.

1. Introduction

Temporal action detection has drawn increasing attention from the research community, owing to its numerous potential applications in surveillance, video analytics, and other areas. This task is to detect human activity instances from untrimmed, and possibly very long videos. Compared to action recognition, it is substantially more challenging, as it is expected to output not only the action category, but also the precise starting and ending time points.

^{*}Indicates equal contributions.

¹Code available at <http://yjsxiong.me/others/ssn>

Over the past several years, the advances in convolutional neural networks have led to remarkable progress in video analysis. Notably, the accuracy of action recognition has been significantly improved [39, 44, 10, 49, 50]. Yet, the performances of action detection methods remain unsatisfactory (see Fig. 1). A major challenge of this problem arises from the large number of highly distinctive but incomplete activity fragments produced by the proposal methods. From the perspective of a snippet-based visual classifier, such incomplete segments are difficult to be distinguished from complete instances. However, to precisely localize an action instance, this distinction between *incomplete* and *complete* segments is crucial. Tackling this chal-

lenge requires the capability of temporal structure analysis, or in other words, the ability to identify different stages *e.g. starting, course, and ending*, which together decide the *completeness* of an activity instance.

Structural analysis is not new in computer vision. It has been well studied in various tasks, *e.g.* image segmentation [21], scene understanding [17], and human pose estimation [2]. Take the most related object detection for example, in deformable part based models (DPM) [9], the modeling of the spatial configurations among parts is crucial. Even with the strong expressive power of convolutional networks [13], explicitly modeling spatial structures, in the form of spatial pyramids [23, 15], remains an effective way to achieve improved performance, as demonstrated in a number of state-of-the-art object detection frameworks, *e.g.* Fast R-CNN [12] and region-based FCN [24].

In the context of video understanding, although temporal structures have played an crucial roles in action recognition [29, 48, 32, 51], its modeling in temporal action detection was not as common and successful. Snippet based methods [25, 41] often process individual snippets independently without considering the temporal structures among them. Later works attempt to incorporate temporal structures, but are often limited to analyzing short clips. S-CNN [37] models the temporal structures via the 3D convolution, but its capability is restricted by the underlying architecture [44], which is designed to accommodate only 16 frames. The methods based on recurrent networks [5, 27] rely on dense snippet sampling and thus are confronted with serious computational challenges when modeling long-term structures. Overall, existing works are limited in two key aspects. First, the tremendous amount of visual data in videos restricts their capability of modeling long-term dependencies in an end-to-end manner. Also, they neither provide *explicit* modeling of different stages in an activity (*e.g. starting and ending*) nor offer a mechanism to assess the *completeness*, which, as mentioned, is crucial for accurate action detection.

In this work, we aim to move beyond these limitations and develop an effective technique for temporal action detection. Specifically, we adopt the proven paradigm of “proposal+classification”, but take a significant step forward, extending the success of explicit structural modeling from the spatial domain to the temporal dimension. In our model, each complete activity instance is considered as a composition of three major stages, namely *starting, course, and ending*, where the stages of starting and ending correspond to the context of an activity instance. Following this notion, we build a temporal pyramid over each proposed segment, which, by combining the features pooled from different stages, will produce a global representation of the entire proposal. Upon the pyramid, we introduce two discriminative models respectively for predicting action cate-

gories and determining *completeness*. Their combination can classify proposed segments to three types: *positive instances, pure backgrounds, and poorly localized proposals*. The temporal pyramid and these discriminative models are integrated into a unified network, called *structured segment network* (SSN). Moreover, we adopt the sparse snippet sampling strategy [50], which overcomes the computational issue for long-term modeling and enables efficient end-to-end training over long videos. Additionally, we propose to use multi-scale grouping upon the temporal actionness signal to generate action proposals. This new method achieves higher temporal recall with less proposals and helps to further boost the detection performance.

The major contribution of this work is the proposal of the structured segment network, a novel framework for temporal action detection that is distinguished in the following aspects: 1) It provides an effective mechanism to model the temporal structures of activities, and thus the capability of discriminating between complete and incomplete proposals. 2) It can be efficiently learned in an end-to-end fashion (5 to 15 hours over a large video dataset, *e.g.* ActivityNet), and once trained, can perform fast inference of temporal structures. 3) The method achieves superior detection performance on standard benchmark datasets. Specifically, it obtains 28.28% average mAP on the ActivityNet [8] v1.3 dataset, outperforming the previous state-of-the-art by 10%.

2. Related Work

Action Recognition. Action recognition has been extensively studied in the past few years [22, 46, 39, 44, 49, 50, 56]. Earlier methods are mostly based on hand-crafted visual features [22, 46]. In the past several years, the wide adoption of convolutional networks (CNNs) has resulted in remarkable performance gain. CNNs are first introduced to this task in [20]. Later, two-stream architectures [39] and 3D-CNN [44] are proposed to incorporate both appearance and motion features. These methods are primarily frame-based and snippet-based, with simple schemes to aggregate results. There are also efforts that explore long-range temporal structures via temporal pooling or RNNs [49, 28, 5]. However, most methods assume well-trimmed videos, where the action of interest lasts for nearly the entire duration. Hence, they don’t need to consider the issue of localizing the action instances.

Object Detection. Our action detection framework is closely related to object detection frameworks [9, 13, 34] in spatial images, where detection is performed by classifying object proposals into foreground classes and a background class. Traditional object proposal methods rely on dense sliding windows [9] and bottom-up methods that exploit low-level boundary cues [45, 57]. Recent proposal methods based on deep neural networks show better average re-

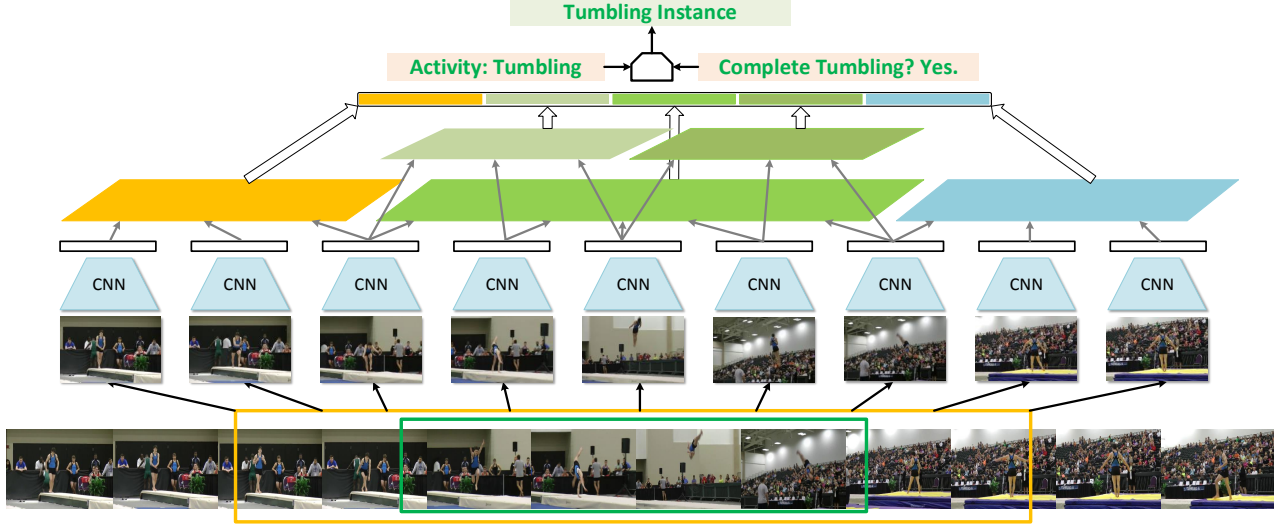


Figure 2. An overview of the structured segment network framework. On a video from ActivityNet [8] there is a candidate region (green box). We first build the augmented proposal (yellow box) by extending it. The augmented proposal is divided into starting (orange), course (green), and ending (blue) stages. An additional level of pyramid with two sub-parts is constructed on the course stage. Features from CNNs are pooled within these five parts and concatenated to form the global region representations. The activity classifier and the completeness classifier operate on the the region representations to produce activity probability and the completeness probability. The final probability of the proposal being positive instance is decided by the joint probability from these two classifiers. During training, we sparsely sample $L = 9$ snippets from evenly divided segments to approximate the dense temporal pyramid pooling.

call while requiring less candidates [34]. Deep models also introduces great modeling capacity for capturing object appearances. With strong visual features, spatial structural modeling [23] remains a key component for detection. In particular, the RoI pooling [12] is introduced to model the spatial configuration of object with minimal extra cost. The idea is further reflected in R-FCN [24] where the spatial configuration is handled with the position sensitive pooling.

Temporal Action Detection. Previous works on activity detection mainly use sliding windows as candidates and focus on designing hand-crafted feature representations for classification [11, 43, 30, 25, 55, 18]. Recent works incorporate deep networks into the detection frameworks and obtain improved performance [54, 37, 3]. S-CNN [37] proposes a multi-stage CNN which boosts accuracy via a localization network. However, S-CNN relies on C3D [44] as the feature extractor, which is initially designed for snippet-wise action classification. Extending it to detection with possibly long action proposals needs enforcing an undesired large temporal kernel stride. Another work [54] uses Recurrent Neural Network (RNN) to learn a glimpse policy for predicting the starting and ending points of an action. Such sequential prediction is often time-consuming for processing long videos and it does not support joint training of the frame-wise CNN for feature extraction. Our method differs from these approaches in that it explicitly models the action structure via structural temporal pyramid pooling. By using sparse sampling, we further enable efficient end-to-end

training. Note there are also works on spatial-temporal detection [14, 53, 26] and temporal video segmentation [16], which are beyond the scope of this paper.

3. Structured Segment Network

The proposed structured segment network framework, as shown in Figure 2, takes as input a video and a set of temporal action proposals. It outputs a set of predicted *activity instances* each associated with a category label and a temporal range (delimited by a starting point and an ending point). From the input to the output, it takes three key steps. First, the framework relies on a proposal method to produce a set of *temporal proposals* of varying durations, where each proposal comes with a starting and an ending time. The proposal methods will be discussed in detail in Section 5. Upon each proposal, a temporal pyramid will then be constructed, which would yield a global representation of the proposal. Finally, two classifiers respectively for recognizing the activity category and assessing the completeness will be applied and their predictions will be combined, resulting in a subset of instances tagged with category labels. Other proposals, which are considered as either *belonging to background* or *incomplete*, will be filtered out.

In this section, we focus on how predictions are made on given proposals. Specifically, our framework considers each proposal as a composition of three consecutive stages, *starting*, *course*, and *ending*, which respectively capture how the action starts, proceeds, and ends. These three

stages form the basis of our structural modeling. Given a proposal, we first divide it into these three stages, and build a separated representation for each stage based on snippet-wise appearance and motion features. The stage-wise representations are integrated via *temporal pyramid pooling* into a *global representation* of the entire proposal. Thereon, we train two classifiers, one for classifying the proposals into C classes plus a “background” class and the other for determining the proposal that captures a *complete* activity. All the components outlined above are integrated into a unified network, which will be trained in an end-to-end way. For training, we adopt the sparse snippet sampling strategy [50] to approximate the temporal pyramid on dense samples. By exploiting the redundancy among video snippets, this strategy can substantially reduce the computational cost, thus allowing the modeling of long-term temporal structures, which is important for many activity categories.

3.1. Three-Stage Structures

At the input level, a video can be represented as a sequence of T snippets, denoted as $(S_t)_{t=1}^T$. Here, one snippet contains several consecutive frames, which, as a whole, is characterized by a combination of RGB images and an optical flow stack [39]. Consider a given set of N proposals $P = \{p_i = [s_i, e_i]\}_{i=1}^N$. Each proposal p_i is composed of a starting time s_i and an ending time e_i . The duration of p_i is thus $d_i = e_i - s_i$. To allow structural analysis and particularly to determine whether a proposal captures a *complete* instance, we need to put it in a context. Hence, we augment each proposal p_i into $p'_i = [s'_i, e'_i]$ with where $s'_i = s_i - d_i/2$ and $e'_i = e_i + d_i/2$. In other words, the augmented proposal p'_i doubles the span of p_i by extending beyond the starting and ending points, respectively by $d_i/2$. If a proposal accurately aligns well with a completed instance, the augmented proposal will capture not only the inherent process of the activity, but also how it starts and how it ends. Following the three-stage notion, we divide the augmented proposal p'_i into three consecutive intervals: $p_i^s = [s'_i, s_i]$, $p_i^c = [s_i, e_i]$, and $p_i^e = [e_i, e'_i]$, which are respectively corresponding to the *starting*, *course*, and *ending* stages.

3.2. Structured Temporal Pyramid Pooling

As mentioned, the structured segment network framework derives a global representation for each proposal via temporal pyramid pooling. This design is inspired by the success of spatial pyramid pooling [23, 15] in object recognition and scene classification. Specifically, given an augmented proposal p'_i divided into three stages p_i^s , p_i^c , and p_i^e , we first compute the stage-wise feature vectors \mathbf{f}_i^s , \mathbf{f}_i^c , and \mathbf{f}_i^e respectively via temporal pyramid pooling, and then concatenate them into a global representation.

Specifically, a stage with interval $[s, e]$ would cover a series of snippets, denoted as $\{S_t | s \leq t \leq e\}$. For each

snippet, we can obtain a feature vector \mathbf{v}_t . Note that we can use any feature extractor here. In this work, we adopt an effective feature representation first proposed in [39]. Based on these features, we construct a K -level temporal pyramid as follows, each level evenly dividing the interval into B_k parts. For the i -th part of the k -th level, whose interval is $[s_{ki}, e_{ki}]$, we can derive a pooled feature as

$$\mathbf{u}_i^{(k)} = \frac{1}{|e_{ki} - s_{ki} + 1|} \sum_{t=s_{ki}}^{e_{ki}} \mathbf{v}_t. \quad (1)$$

Then the overall representation of this stage can be obtained by concatenating the pooled features across all parts at all levels as $\mathbf{f}_i^c = (\mathbf{u}_i^{(k)})_{k=1, \dots, K, i=1, \dots, B_k}$.

We treat the three stages differently. Generally, we observed that the *course* stage, which reflects the activity process itself, usually contains richer structure *e.g.* this process itself may contain sub-stages. Hence, we use a two-stage pyramid, *i.e.* $K = 2$, for the *course* stage, while using simpler one-stage pyramids (which essentially reduce to standard average pooling) for *starting* and *ending* pyramids. We found empirically that this setting strikes a good balance between expressive power and complexity. Finally, the stage-wise features are combined via concatenation. Overall, this construction explicitly leverages the structure of an activity instance and its surrounding context, and thus we call it *structured temporal pyramid pooling*.

3.3. Activity and Completeness Classifiers

On top of the structured features described above, we introduce two types of classifiers, an *activity classifier* and a set of *completeness classifiers*. Specifically, the *activity classifier* A classifies input proposals into $K + 1$ classes, *i.e.* K activity classes (with labels $1, \dots, K$) and an additional “background” class (with label 0). This classifier restricts its scope to the *course* stage, making predictions based on the corresponding feature \mathbf{f}_i^c . The *completeness classifiers* $\{C_k\}_{k=1}^K$ are a set of binary classifiers, each for one activity class. Particularly, C_k predicts whether a proposal captures a *complete* activity instance of class k , based on the global representation $\{\mathbf{f}_i^s, \mathbf{f}_i^c, \mathbf{f}_i^e\}$ that covers all three stages. In this way, the *completeness* is determined not only on the proposal itself but also on its surrounding context.

Both types of classifiers are implemented as linear classifiers on top of high-level features. Given a proposal p_i , the activity classifier will produce a vector of normalized responses via a softmax layer. From a probabilistic view, it can be considered as a conditional distribution $P(c_i | p_i)$, where c_i is the class label. For each activity class k , the corresponding completeness classifier C_k will yield a probability value, which can be understood as the conditional probability $P(b_i | c_i, p_i)$, where b_i indicates whether p_i is *complete*. Both outputs together form a joint distribution.

When $c_i \geq 1$, $P(c_i, b_i | p_i) = P(c_i | p_i) \cdot P(b_i | c_i, p_i)$. Hence, we can define a *unified classification loss* jointly on both types of classifiers. With a proposal p_i and its label c_i :

$$\mathcal{L}_{cls}(c_i, b_i; p_i) = -\log P(c_i | p_i) - 1_{(c_i \geq 1)} P(b_i | c_i, p_i). \quad (2)$$

Here, the *completeness* term $P(b_i | c_i, p_i)$ is only used when $c_i \geq 1$, i.e. the proposal p_i is not considered as part of the background. Note that these classifiers together with the structured temporal pyramid are integrated into a single network that is trained in an *end-to-end* way.

During training, we collect three types of proposal samples: (1) *positive proposals*, i.e. those overlap with the closest groundtruth instances with at least 0.7 IoU; (2) *background proposals*, i.e. those that do not overlap with any groundtruth instances; and (3) *incomplete proposals*, i.e. those that satisfy the following criteria: 80% of its own span is contained in a groundtruth instance, while its IoU with that instance is below 0.3 (in other words, it just covers a small part of the instance). For these proposal types, we respectively have $(c_i > 0, b_i = 1)$, $c_i = 0$, and $(c_i > 0, b_i = 0)$. Each mini-batch is ensured to contain all three types of proposals.

3.4. Location Regression and Multi-Task Loss

With the structured information encoded in the global features, we can not only make categorical predictions, but also refine the proposal’s temporal interval itself by *location regression*. We devise a set of location regressors $\{R_k\}_{k=1}^K$, each for an activity class. We follow the design in RCNN [13], but adapting it for 1D temporal regions. Particularly, for a *positive proposal* p_i , they regress both the interval center μ_i and the span ϕ_i (in log-scale), using the closest groundtruth instance as the target. With both the classifiers and location regressors, we define a multi-task loss over an training sample p_i , as:

$$\mathcal{L}_{cls}(c_i, b_i; p_i) + \lambda \cdot 1_{(c_i \geq 1 \ \& \ b_i = 1)} \mathcal{L}_{reg}(\mu_i, \phi_i; p_i). \quad (3)$$

Here, \mathcal{L}_{reg} uses the smooth L_1 loss function [12].

4. Efficient Training and Inference with SSN

The huge computational cost poses a serious challenge to video analysis. Our structured segment network framework also faces this challenge. This section presents two techniques which we use to reduce the cost.

Training with sparse sampling. The structured temporal pyramid, in its original form, rely on densely sampled snippets. This would lead to excessive computational cost and memory demand when training over long proposals – in practice, proposals that span over hundreds of frames are not uncommon. However, dense sampling is generally unnecessary in our framework. Particularly, the *pooling* operation is essentially to collect feature statistics over a certain

region. Such statistics can be well approximated via a subset of snippets, due to the high redundancy among them.

Motivated by this, we devise a *sparse snippet sampling scheme*. Specifically, given a augmented proposal p'_i , we evenly divide it into $L = 9$ segments, randomly sampling only one snippet from each segment. Structured temporal pyramid pooling is performed for each pooling region on its corresponding segments. This scheme is inspired by the segmental architecture in [50], but differs in that it operates with the pyramids instead of a global average pooling. In this way, we fix the number of features needed to be computed regardless of how long the proposal is, thus effectively reducing the computational cost, especially for modeling long-term structures. More importantly, this enables end-to-end training of the entire framework over a large number of long proposals.

Inference with reordered computation. In testing, we sample video snippets with a fixed interval of 6 frames, and construct the temporal pyramid thereon. The original formulation of temporal pyramid first computes pooled features and then applies the classifiers and regressors on top. This, however, is not a very efficient way. Actually, for each video, hundreds of proposals will be generated, and these proposals can significantly overlap with each other – therefore, a considerable portion of the snippets and thus the features derived thereon are shared among proposals.

To exploit this redundancy in the computation, we adopt the idea introduced in position sensitive pooling [24] to improve testing efficiency. Note that our classifiers and regressors are both linear. So the key step in classification or regression is to multiply a weight matrix \mathbf{W} with the global feature vector \mathbf{f} . Recall that \mathbf{f} itself is a concatenation of multiple features, each pooled over a certain interval. Hence the computation can be written as $\mathbf{W}\mathbf{f} = \sum_j \mathbf{W}_j \mathbf{f}_j$, where j indexes different regions along the pyramid. Here, \mathbf{f}_j is obtained by *average pooling* over all snippet-wise features within the region r_j . Thus, we have

$$\mathbf{W}_j \mathbf{f}_j = \mathbf{W}_j \cdot \mathbb{E}_{t \sim r_j} [\mathbf{v}_t] = \mathbb{E}_{t \sim r_j} [\mathbf{W}_j \mathbf{v}_t]. \quad (4)$$

$\mathbb{E}_{t \sim r_j}$ denotes the average pooling over r_j , which is a linear operation and therefore can be exchanged with the matrix multiplication. Eq (4) suggests that the linear responses w.r.t. the classifiers/regressors can be computed *before* pooling. In this way, the heavy matrix multiplication can be done once for each video over all snippets, and for each proposal, we only have to pool over the response values. After extracting the network outputs, by reordering matrix multiplication and pooling, the inference time is reduced from around 10 seconds to less than 0.5 second per video on average.




Figure 3. Visualization of the temporal actionness grouping process for proposal generation. **Top:** Actionness probabilities as a 1D signal sequence. **Middle:** The complement signal. We flood it with different levels γ . **Bottom:** Regions obtained by different flooding levels. By merging the regions according to the grouping criterion, we get the final set of proposals (in orange color).

5. Temporal Region Proposals

Our framework accepts arbitrary proposals, *e.g.* sliding windows [37, 55]. Yet, an effective proposal method can produce more accurate proposals, and thus allowing a small number of proposals to reach a certain level of performance. In this work, we devise a simple yet effective proposal method called *temporal actionness grouping (TAG)*.

This method uses an actionness classifier to evaluate the *actionness probabilities* for individual snippets. The actionness classifier is a binary classifier that can be trained on snippets. Our basic idea is to find those continuous regions with mostly high actionness probabilities. To this end, we repurpose a classic watershed algorithm [36], applying it to an 1D signal formed by a sequence of *complemented* actionness values ($1 - a$), as shown in Figure 3. This algorithm floods the surface with different “water level” (γ), resulting in a set of “basins”, denoted by $G(\gamma)$. Intuitively, each “basin” corresponds to a temporal region with high actionness.

Given a set of basins $G(\gamma)$, we devise a grouping scheme similar to [33], which tries to connect small basins into proposal regions. The scheme works as follows: it begins with a seed basin, and consecutively absorbs the basins that follow, until the fraction of the basin parts over the total duration (*i.e.* from the beginning of the first basin to the ending of the last) drops below a certain threshold γ . In this way, it can produce a set of regions starting from different seeds, denoted by $G'(\tau, \gamma)$. Note that we do not choose a specific combination of τ and γ . Instead we uniformly sample τ and γ from $\in (0, 1)$ with an even step of 0.05. The com-

bination of these two thresholds leads to multiple sets of regions. We then take the *union* of them. Finally, we apply non-maximal suppression to the union with IoU threshold 0.95, to filter out highly overlapped proposals. The retained proposals will be fed to the SSN framework.

6. Experimental Results

We conducted experiments to test the proposed framework on two large-scale action detection benchmark datasets: *ActivityNet* [8] and *THUMOS14* [19]. This section first introduces these datasets and other experimental settings, then compares the performance with other state-of-the-art approaches, and finally investigate the impact of different components via a set of ablation studies.

6.1. Experimental Settings

Datasets *ActivityNet* [8] has two versions, *v1.2* and *v1.3*. The former contains 9682 videos in 100 classes, while the latter, which is a superset of *v1.2* and was used in the *ActivityNet Challenge 2016*, contains 19994 videos in 200 classes. In each version, the dataset is divided into three disjoint subsets, training, validation, and testing, by 2:1:1. *THUMOS14* [19] has 1010 videos for validation and 1574 videos for testing. These videos contain annotated action instances in 20 classes. This dataset does not provide the training set by itself. Instead, the UCF101 [42] is appointed as the official training set. As no temporal annotations are provided by the training set, we instead train the models on the validation set and evaluate them on the testing set. So there are about 220 videos from 20 action classes for training. In our experiments, we compare with our method with the states of the art on both *THUMOS14* and *ActivityNet v1.3*, and perform ablation studies on *ActivityNet v1.2*.

Implementation Details We train the structured segment network in an end-to-end manner, with raw video frames and action proposals as the input. Two-stream CNNs [39] are used for feature extraction. We also use the spatial and temporal streams to harness both the appearance and motion features. The binary actionness classifiers underlying the TAG proposals are trained with [50] on the training subset of each dataset. We use SGD to learn CNN parameters in our framework, with batch size 128 and momentum 0.9. We initialize the CNNs with pre-trained models from ImageNet [4]. The initial learning rates are set to 0.001 for RGB networks and 0.005 for optical flow networks. In each minibatch, we keep the ratio of three types of proposals, namely *positive*, *background*, and *incomplete*, to be 1:1:6. For the completeness classifiers, only the samples with loss values ranked in the first 1/6 of a minibatch are used for calculating gradients, which resembles online hard negative mining [38]. On both versions of *ActivityNet*, the RGB and optical flow branches of the two-stream CNN are respectively trained for 9.5K and 20K iterations, with learning




Figure 4. Recall rate at different tIoU thresholds on ActivityNet v1.2. A high recall at high IoU threshold (> 0.7) usually reflects better proposal quality.

Proposal Method	THUMOS14		ActivityNet v1.2	
	# Prop.	AR	# Prop.	AR
Sliding Windows	204	21.2	100	34.8
SCNN-prop [37]	200	20.0	-	-
TAP [7]	200	23.0	90	14.9
DAP [6]	200	37.0	100	12.1
TAG	200	48.9	100	67.3

Table 1. Comparison between different temporal action proposal methods. “AR” refers to the average recall rates. “-” indicates the result is not available. We use roughly same numbers of proposals for each method.

rates scaled down by 0.1 after every $4K$ and $1K$ iterations, respectively. On THUMOS14, these two branches are respectively trained for $1K$ and $6K$ iterations, with learning rates scaled down by 0.1 per 400 and 2500 iterations. We will release all the models and the source codes.

Evaluation Metrics As both datasets originate from contests, each dataset has its own convention of reporting performance metrics. We follow their conventions, reporting mean average precision (mAP) at different IoU thresholds. On both versions of ActivityNet, the IoU thresholds are $\{0.5, 0.75, 0.95\}$. The average of mAP values with IoU thresholds $[0.5:0.05:0.95]$ is used to compare the performance between different methods. On THUMOS14, the IoU thresholds are $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. The mAP at 0.5 IoU is used for comparing results from different methods.

6.2. Ablation Studies

Temporal Action Proposal We compare the performance of different action proposal schemes in three aspects, *i.e.*

Average mAP (%)	(1)-0	(1,2)-0	(1,2)-1
Max Pool	13.1	13.5	18.1
Average Pool	4.48	4.34	23.7

Table 2. Comparison between different temporal pooling settings. The setting (1,2)-1 is used in the SSN framework. Please refer to Sec. 6.2 for the definition of these settings.

recall, quality, and detection performance. Particularly, we compare our TAG scheme with common sliding windows as well as other state-of-the-art proposal methods, including SCNN-prop, a proposal networks presented in [37], TAP [7], DAP [6]. For the sliding window scheme, we use 20 exponential scales starting from 0.3 second long and step sizes of 0.4 times of window lengths.

We first evaluate the average recall rates, which are summarized in Table 1. We can see that TAG proposal have higher recall rates with the same number of proposals. Then we investigate the quality of its proposals. We plot the recall rates from different proposal methods at different IoU thresholds in Fig. 4. We can see TAG retains relative high recall at high IoU thresholds, demonstrating that the proposals from TAG are generally more accurate. In experiments we also tried applying the actionness classifier trained on ActivityNet v1.2 directly on THUMOS14. We can still achieve a reasonable average recall of 39.6%, while the one trained on THUMOS14 achieves 48.9% in Table 1. Finally, we evaluate the proposal methods in the context of action detection. The detection mAP values using sliding window proposals and TAG proposals are shown in Table 3. The results confirm that, in most cases, the improved proposals can result in improved detection performance.

Structured Temporal Pyramid Pooling. Here we study the influence of different pooling strategies. We denote one pooling configuration as $(B_1, \dots, B_K) - A$, where K refers to the number of pyramid levels for the course stage and B_1, \dots, B_K the number of regions in each level. $A = 1$ indicates we use augmented proposal and model the starting and ending stage, while $A = 0$ indicates we only use the original proposal (without augmentation). Additionally we compare two within-region pooling methods: average and max pooling. The results are summarized in Table 2. Note for simplicity, we only investigate these configurations in the stage-wise training scenarios. The results suggest that the detection performance is generally better when we use pyramid pooling on the augmented proposals, which corroborate the importance of the structured modeling and the use of contextual information. Based on these results, we fix the configuration to $(1, 2) - 1$ in later experiments.

Classifier Design For proposal classification, we have both the activity classifiers and the completeness classifiers. A major distinction between them is that they consider different kinds of samples as “negative”. We verify the impor-

	Stage-Wise				End-to-End	
STPP	✓	✓	✓		✓	✓
Act. + Comp.			✓	✓	✓	✓
Loc. Reg.				✓		✓
SW	0.56	2.26	16.4	18.1	-	-
TAG	4.82	9.55	23.7	24.2	23.8	24.5

Table 3. Study of the contribution of each component in the SSN framework. Overall, end-to-end training is compared against stage wise training. We evaluate the performance using both sliding window proposals (“SW”) and TAG proposals (“TAG”). Performance is measured by average mAP on ActivityNet [8] v1.2. In this table, “STPP” refers to structure temporal pyramid pooling. “Act. + Comp.” refers to the use of two classifiers design. “Loc. Reg.” denotes the use the location regression.

tance of this decomposed design by studying another design that replaces it with a single set of classifiers, for which both *background* and *incomplete* samples are uniformly treated as negative. The results are also summarized in Table 3. We observe that using only one classifier to distinguish *positive* samples from both *background* and *incomplete* would lead to substantially worse results, with mAP decreased from 23.7% to 9.55%. The performance gain of the decomposed classifier design is due to the different natures of the two negative proposal types, which require different classifiers to handle.

Location Regression & Multi-Task Learning Because of the contextual information contained in the starting and ending stages of the global region features, we are able to regress to the possibly true location of a proposed region. We measure the contribution of location regression to the detection performance in Table 3. It is also important to evaluate the effectiveness of the multi-task training, where we train the classifiers with the regressors together. It is used in the end-to-end training scenarios. From the results we can see that the location regression always improves the detection accuracy, and the multi-task learning can lead to even better performance.

Training: Stage-wise v.s. End-to-end. While the structured segment network is designed for end-to-end training, it is also possible to first extract features and train the classifiers and regressors with SVM and ridge regression, respectively. We refer to this training scheme as stage-wise training. We compare the performance of end-to-end training and stage-wise training in Table 3. We observe that models from end-to-end training always outperform those learned with stage-wise training under the same settings. This is remarkable as we are only sparsely sampling snippets in the former setting while the latter pools features on densely sampled snippets. It clearly demonstrates the importance of jointly optimizing the classifiers and feature extractors and justifies our framework design.

THUMOS14, mAP@ α					
Method	0.1	0.2	0.3	0.4	0.5
Wang <i>et. al.</i> [47]	18.2	17.0	14.0	11.7	8.3
Oneata <i>et. al.</i> [31]	36.6	33.6	27.0	20.8	14.4
Richard <i>et. al.</i> [35]	39.7	35.7	30.0	23.2	15.2
S-CNN [37]	47.7	43.5	36.3	28.7	19.0
Yeung <i>et. al.</i> [54]	48.9	44.0	36.0	26.4	17.1
Yuan <i>et. al.</i> [55]	51.4	42.6	33.6	26.1	18.8
SSN	66.0	59.4	51.9	41.0	29.8

Table 4. Action detection results on THUMOS14, measured by mean average precision (mAP) for different IoU thresholds α . The upper half of the table shows challenge results back in 2014.

ActivityNet v1.3 (testing), mAP@ α				
Method	0.5	0.75	0.95	Average
Wang <i>et. al.</i> [52]	42.48	2.88	0.06	14.62
Singh <i>et. al.</i> [40]	28.67	17.78	2.88	17.68
Singh <i>et. al.</i> [41]	36.40	11.05	0.14	17.83
SSN	43.26	28.70	5.63	28.28

Table 5. Action detection results on ActivityNet v1.3, measured by mean average precision (mAP) for different IoU thresholds α and the average mAP of IoU thresholds from 0.5 to 0.95.

6.3. Comparison with the State of the Art

Finally, we compare our method with other state-of-the-art temporal action detection methods on THUMOS14 [19] and ActivityNet v1.3 [8], and report the performances using the metrics described above. Note that the average action duration in THUMOS14 is 4 seconds, while those in ActivityNet datasets can be as long as 50 seconds. Meanwhile the average video duration in THUMOS14 is about twice than those of the ActivityNet datasets (233 sec v.s. 114 sec). This difference in the statistics reflects the different natures of these datasets in terms of the granularities and temporal structures of the action instances. Hence, strong adaptivity is required to perform consistently well on both datasets.

THUMOS14 On THUMOS 14, We compare with the contest results [47, 31, 35] and those from recent works, including the methods that use segment-based 3D CNN [37], score pyramids [55], and recurrent reinforcement learning [54]. The results are shown in Table 4. In most cases, the proposed method outperforms previous state-of-the-art methods by over 10% in absolute mAP values.

ActivityNet The results on the testing set of ActivityNet v1.3 are shown in Table 5. For references, we list the performances of highest ranked entries in the ActivityNet 2016 challenge. We submit our results to the test server of ActivityNet v1.3 and report the detection performance on the testing set. The proposed framework, using a single model instead of an ensemble, is able to outperform the winning

submission at every IoU threshold and achieves an average mAP of 28.28. Under higher IoU thresholds, *i.e.*, 0.75 and 0.95, the proposed detector performs drastically better than all rivals, *e.g.* with IoU threshold 0.75, we obtain an mAP at 28.7% (vs 17.8%); with IoU threshold 0.95, we obtain an mAP at 5.63% (vs 2.88%). In terms of average mAP, we are able to outperform other state-of-the-arts by around 10%. This clearly demonstrates the superiority of our method. Visualization of the detection results can be found in the supplementary materials [1].

7. Conclusion

In this paper, we presented a generic framework for temporal action detection, which combines a structured temporal pyramid with two types of classifiers, respectively for predicting activity class and completeness. With this framework, we achieved significant performance gain over state-of-the-art methods on both ActivityNet and THUMOS14. Moreover, we demonstrated that our method is both accurate and generic, being able to localize temporal boundaries precisely and working well for activity classes with very different temporal structures.

References

- [1] Supplementary materials. <http://yjxiong.me/others/ssn/supp.pdf>. 9
- [2] M. Andriluka, S. Roth, and B. Schiele. Pictorial structures revisited: People detection and articulated pose estimation. In *CVPR*, pages 1014–1021. IEEE, 2009. 2
- [3] R. De Geest, E. Gavves, A. Ghodrati, Z. Li, C. Snoek, and T. Tuytelaars. Online action detection. In *ECCV*, pages 269–284. Springer, 2016. 3
- [4] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 6
- [5] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pages 2625–2634, 2015. 2
- [6] V. Escorcia, F. Caba Heilbron, J. C. Niebles, and B. Ghanem. Daps: Deep action proposals for action understanding. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *ECCV*, pages 768–784, 2016. 7
- [7] B. G. Fabian Caba Heilbron, Juan Carlos Niebles. Fast temporal activity proposals for efficient detection of human actions in untrimmed videos. In *CVPR*, pages 1914–1923, 2016. 7
- [8] B. G. Fabian Caba Heilbron, Victor Escorcia and J. C. Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. 2, 3, 6, 8
- [9] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE TPAMI*, 32(9):1627–1645, 2010. 2
- [10] B. Fernando, E. Gavves, J. O. M., A. Ghodrati, and T. Tuytelaars. Modeling video evolution for action recognition. In *CVPR*, pages 5378–5387, 2015. 1
- [11] A. Gaidon, Z. Harchaoui, and C. Schmid. Temporal localization of actions with actoms. *IEEE TPAMI*, 35(11):2782–2795, 2013. 3
- [12] R. Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015. 2, 3, 5
- [13] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014. 2, 5
- [14] G. Gkioxari and J. Malik. Finding action tubes. In *CVPR*, pages 759–768, June 2015. 3
- [15] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *ECCV*, pages 346–361. Springer, 2014. 2, 4
- [16] M. Hoai, Z.-Z. Lan, and F. De la Torre. Joint segmentation and classification of human actions in video. In *CVPR*, pages 3265–3272. IEEE, 2011. 3
- [17] D. Hoiem, A. A. Efros, and M. Hebert. Putting objects in perspective. *IJCV*, 80(1):3–15, 2008. 2
- [18] M. Jain, J. C. van Gemert, H. Jégou, P. Bouthemy, and C. G. M. Snoek. Action localization by tubelets from motion. In *CVPR*, June 2014. 3
- [19] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://csrcv.ucf.edu/THUMOS14/>, 2014. 6, 8
- [20] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, pages 1725–1732, 2014. 2
- [21] J. Lafferty, A. McCallum, F. Pereira, et al. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, volume 1, pages 282–289, 2001. 2
- [22] I. Laptev. On space-time interest points. *IJCV*, 64(2-3), 2005. 2
- [23] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, volume 2, pages 2169–2178. IEEE, 2006. 2, 3, 4
- [24] Y. Li, K. He, J. Sun, et al. R-fcn: Object detection via region-based fully convolutional networks. In *NIPS*, pages 379–387, 2016. 2, 3, 5
- [25] P. Mettes, J. C. van Gemert, S. Cappallo, T. Mensink, and C. G. Snoek. Bag-of-fragments: Selecting and encoding video fragments for event detection and recounting. In *ICMR*, pages 427–434, 2015. 2, 3
- [26] P. Mettes, J. C. van Gemert, and C. G. Snoek. Spot on: Action localization from pointly-supervised proposals. In *ECCV*, pages 437–453. Springer, 2016. 3
- [27] A. Montes, A. Salvador, S. Pascual, and X. Giro-i Nieto. Temporal activity detection in untrimmed videos with recurrent neural networks. In *NIPS Workshop*, 2016. 2
- [28] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, pages 4694–4702, 2015. 2

- [29] J. C. Niebles, C.-W. Chen, and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*, pages 392–405. Springer, 2010. 2
- [30] D. Oneata, J. Verbeek, and C. Schmid. Action and event recognition with fisher vectors on a compact feature set. In *ICCV*, pages 1817–1824, 2013. 3
- [31] D. Oneata, J. Verbeek, and C. Schmid. The lear submission at thumos 2014. In *THUMOS Action Recognition Challenge*, 2014. 8
- [32] H. Pirsiavash and D. Ramanan. Parsing videos of actions with segmental grammars. In *CVPR*, pages 612–619, 2014. 2
- [33] J. Pont-Tuset, P. Arbeláez, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. In *arXiv:1503.00848*, March 2015. 6
- [34] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015. 2, 3
- [35] A. Richard and J. Gall. Temporal action detection using a statistical language model. In *CVPR*, pages 3131–3140, 2016. 8
- [36] J. B. Roerdink and A. Meijster. The watershed transform: Definitions, algorithms and parallelization strategies. *Fundamenta informaticae*, 41(1, 2):187–228, 2000. 6
- [37] Z. Shou, D. Wang, and S.-F. Chang. Temporal action localization in untrimmed videos via multi-stage CNNs. In *CVPR*, pages 1049–1058, 2016. 2, 3, 6, 7, 8
- [38] A. Shrivastava, A. Gupta, and R. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, pages 761–769, 2016. 6
- [39] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, pages 568–576, 2014. 1, 2, 4, 6
- [40] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, pages 1961–1970, 2016. 8
- [41] G. Singh and F. Cuzzolin. Untrimmed video classification for activity detection: submission to activitynet challenge. *CoRR*, abs/1607.01979, 2016. 2, 8
- [42] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv:1212.0402*, 2012. 6
- [43] K. Tang, B. Yao, L. Fei-Fei, and D. Koller. Combining the right features for complex event recognition. In *CVPR*, pages 2696–2703, 2013. 3
- [44] D. Tran, L. D. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3D convolutional networks. In *ICCV*, pages 4489–4497, 2015. 1, 2, 3
- [45] K. E. Van de Sande, J. R. Uijlings, T. Gevers, and A. W. Smeulders. Segmentation as selective search for object recognition. In *ICCV*, pages 1879–1886, 2011. 2
- [46] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, pages 3551–3558, 2013. 2
- [47] L. Wang, Y. Qiao, and X. Tang. Action recognition and detection by combining motion and appearance features. In *THUMOS Action Recognition Challenge*, 2014. 8
- [48] L. Wang, Y. Qiao, and X. Tang. Latent hierarchical model of temporal structure for complex activity classification. *IEEE TIP*, 23(2):810–822, 2014. 2
- [49] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *CVPR*, pages 4305–4314, 2015. 1, 2
- [50] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36, 2016. 1, 2, 4, 5, 6
- [51] P. Wang, Y. Cao, C. Shen, L. Liu, and H. T. Shen. Temporal pyramid pooling based convolutional neural network for action recognition. *IEEE TCSVT*, 2016. 2
- [52] R. Wang and D. Tao. UTS at activitynet 2016. In *ActivityNet Large Scale Activity Recognition Challenge 2016*, 2016. 8
- [53] P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Learning to track for spatio-temporal action localization. In *ICCV*, pages 3164–3172, 2015. 3
- [54] S. Yeung, O. Russakovsky, G. Mori, and L. Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, pages 2678–2687, 2016. 3, 8
- [55] J. Yuan, B. Ni, X. Yang, and A. A. Kassim. Temporal action localization with pyramid of score distribution features. In *CVPR*, pages 3093–3102, 2016. 3, 6, 8
- [56] B. Zhang, L. Wang, Z. Wang, Y. Qiao, and H. Wang. Real-time action recognition with enhanced motion vector CNNs. In *CVPR*, pages 2718–2726, 2016. 2
- [57] C. L. Zitnick and P. Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, pages 391–405, 2014. 2